Order Parameters for Visual Inference

Alan Yuille
UCLA
Dept. Statistics and Psychology.

How hard are visual tasks?

Easy, Medium, and Hard target detection.

How hard are visual tasks?

- **What are the factors that determine how hard a task is?**
- When can tasks be solved?

How fast can we solve them?

- What approximations can we make

and still solve them?

Bayes Risk and Performance

- **Bayes Risk. I.**
- I is the input data and W is the

representation we seek to compute.

- Seek a decision rule W=d(I).
- d(I) can be a analytic function, the solution of a PDE, filter response…

Bayes Risk and Performance

- **Bayes Risk. II**.
- Need an error criterion L(W,d(I)) – the loss of making decision d(I) when the true solution is W.
- Need a dataset of problem examples, ***Ensemble of problem instances P(W,I).***

Bayes Risk and Performance

- **Bayes Risk. III.**

$$\text{Risk}: R(d) = \sum_{I,W} L(W, d(I))\, P(W, I).$$

$$\text{Bayes Risk} = \min_{d} R(d).$$

If we restrict the class of decision rules to Ω, then performance degrades to $\min_{d \in \Omega} R(d)$.

Bayesian Decision Theory

- Most existing work on performance analysis of visual algorithms is, implicitly, or explicitly, formulated in these terms.
- Frequentist bounds, ROC curves, Cramer-Rao, Hilbert-Schmidt, etc.
- But there are limited analytic studies for realistic vision problems.

Order Parameters.

- Use Bayesian Decision Theory to obtain analytic performance bounds for road detection.
- Performance will be determined by **order parameters**.
- Phase transitions in performance at critical values of order parameters.

Road Detection

- Road Detection from Aerial Images. (Geman ’96)
- Task: find and track the road.
- Performance and Complexity analysis (Yuille, Coughlan),(Yuille,Coughlan, Wu, Zhu).

Geman and Jedynak Model.

- A path is a sequence of moves $\{t_i\}$ on a Q-nary tree. This determines a path in the image $\{x_i\}$:

$$x_{(i+1)} = x_{(i)} + w(x_{(i)}-x_{(i-1)}:t_{(i)}),$$

- where w(.) is an arc of fixed length and depends on the angle of $t_{(i)}$ relative to Previous segment $x_{(i)}-x_{(i-1)}$.

Geman and Jedynak Tree

- 3-nary tree.

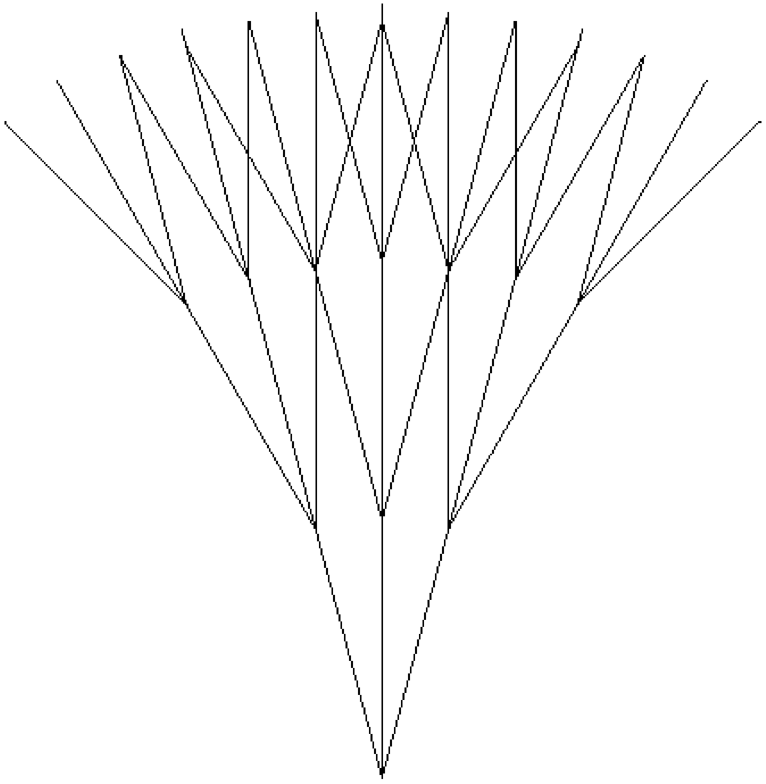

Geman and Jedynak Model

- Geometric prior on path expressed by $P(\{t_{(i)}\}) = P_g(t_{(1)}) P_g(t_{(n)})$ (can generalize to first order Markov).
- Likelihood function: filter response is drawn from $P(y_{(i)}|on)$ or $P(y_{(i)}|off)$ if filter is evaluated on or off the road.

Geometric Prior.

Example of a geometric prior.
Samples of curves with elastica model (Mumford 1995).

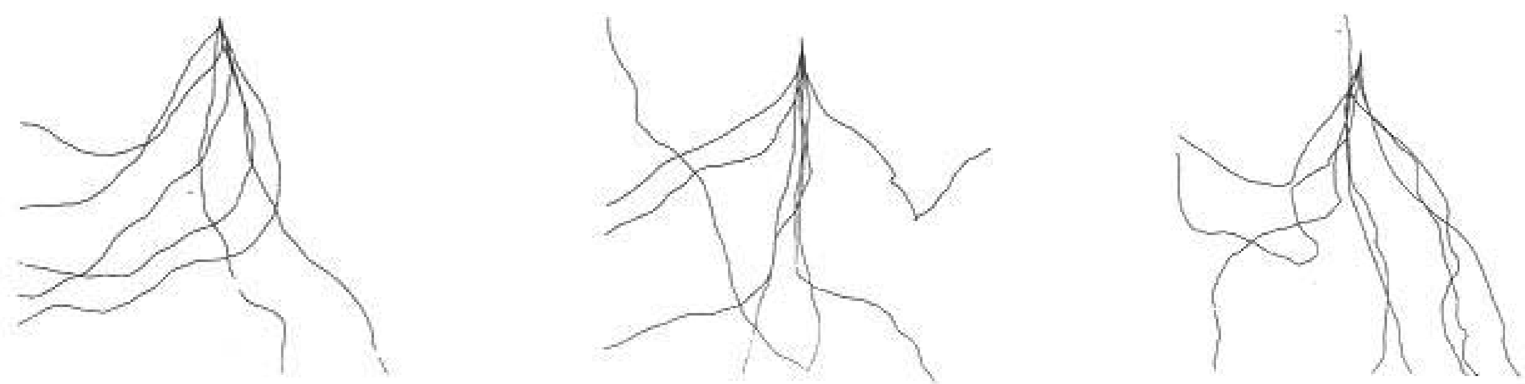

Likelihood Function

- P-on and P-off (empirical).

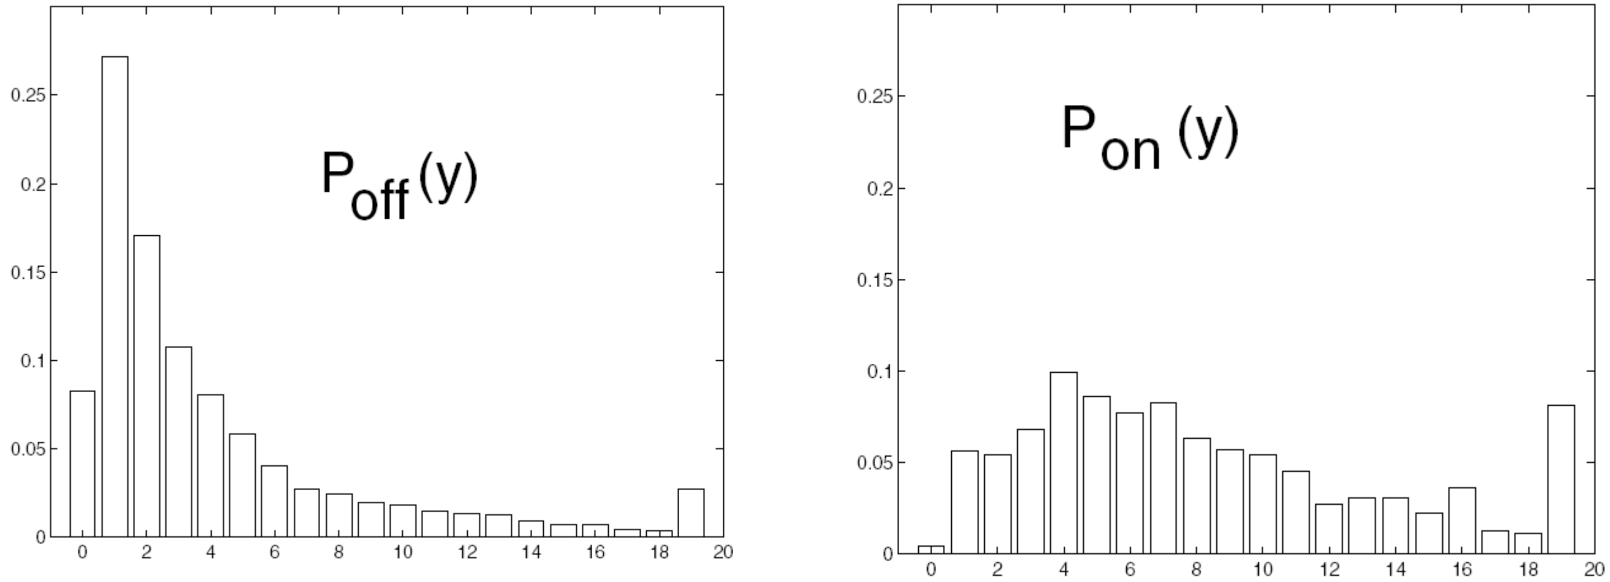

Geman and Jedynak Model

- MAP estimation corresponds to finding the path that maximizes the reward function:

$$r(\{t_i\},\{y_i\}) = \frac{1}{N}\sum_{i=1}^{N} \log \frac{P_{on}(y_i)}{P_{off}(y_i)}$$

$$+\frac{1}{N}\sum_{i=1}^{N} \log \frac{P_g(t_i)}{U(t_i)},$$

$U(.)$ is uniform distribution.

A* Search to get MAP

- Tree Search (Left). A* heuristic.

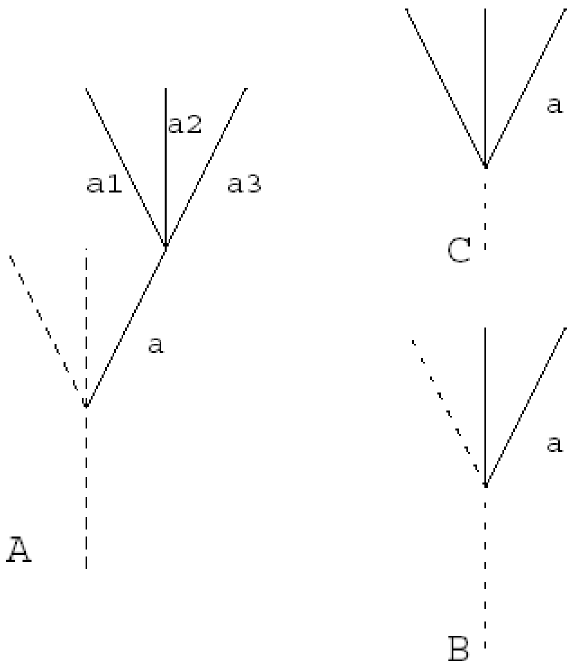

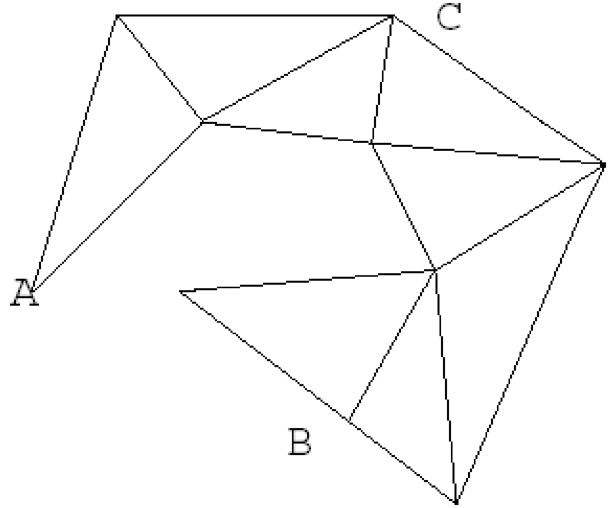

Model Performance

- Geman and Jedynak report good performance with convergence linear in length N of road.
- But, others report (privately) poor results using this algorithm.

Theoretical Analysis

- **Problem formulation defines an ensemble of problem instances.**
- $P(\{t_{(i)}\},\{y_{(i)}\})$ determined by the geometric prior $P(\{t_{(i)}\})$ and the likelihood function $P(\{y_{(i)}\}|\{t_{(i)}\})$.
- Likelihood function is factorized in terms of P_{on} and P_{off}.

Distribution for rewards.

- Calculate the probability distribution for the reward ***on the road path***, with respect to ensemble $P(\{t_{(i)}\},\{y_{(i)}\})$.
- Calculate the distribution for the reward for the Q^N-1 ***off road paths.***
- These calculations are possible because of shift-invariance (self-averaging).

Analysis using Sanov's Thm.

- **Theorem (Phase Transition).**

Let K=D(P_on||P_off) + D(P_g||U) – log Q. Then:

(I) Bayes risk tends to 0 as N tends to infinity if, and only if, K >0.

(II) Bayes risk tends to 1 as N tends to infinity if K<0.

D(.||.) is Kullback-Leibler divergence.

Examples:

- Order Parameter K:
- K=0.8647, K =0.2105, K =-0.7272.

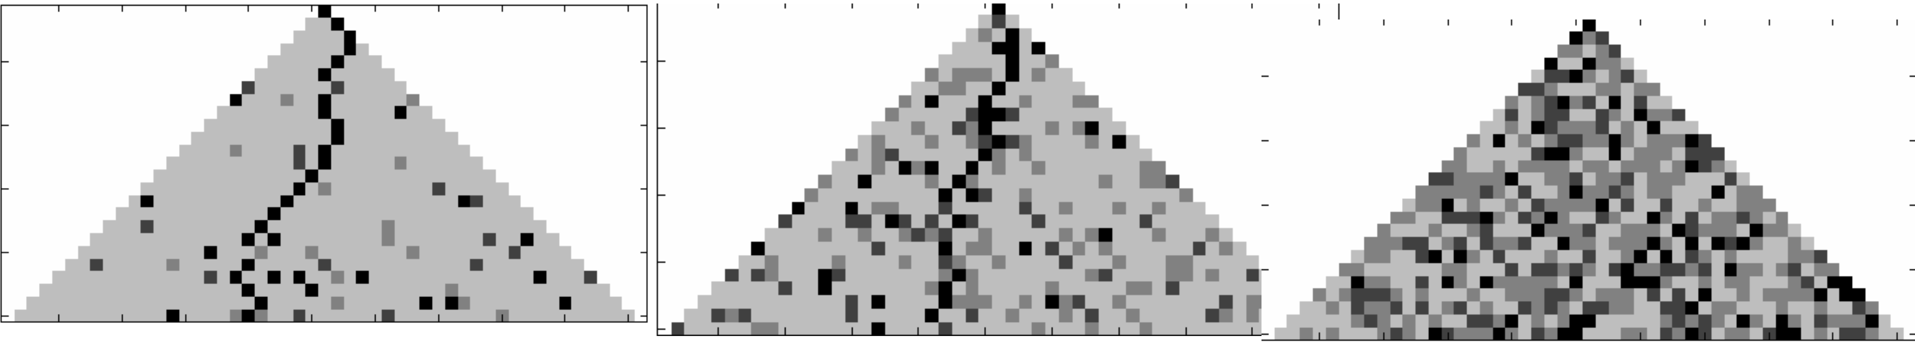

Bounds from Sanov's Thm.

- Properties of interest – error rates – are bounded by terms such as e^ (-N K) for large N.
- Sanov's Theorem – see Cover and Thomas (1991) – applies to i.i.d. samples (requires distributions to be factorized).

Convergence Rate.

- Similar analysis shows:

Expected Complexity is O(N),

if K > T >0, (T is a small constant).

Coefficient of N is bounded by simple function of K.

Worst case is exponential in N. But typical performance is far better.

Ensemble Analysis

- Our analysis shows that the ability to detect the road depends on an order parameter K.

K is a function of the probability distributions describing the problem. (France, Portugal,…).

(I). If K <0, the task cannot be solved by any algorithm (needle in haystack).

(II). If K >0, the task can be solved.

Approximations.

- Perform Inference using an approximate model P_g^a instead of P_g. (sufficient statistics).
- This reduces the order parameter: $K^a = K - D(P\_g||P\_g^a)$.
- This reduces performance by a quantifiable amount -- $D(P\_g||P\_g^a)$.

More General Models

- Generalize previous results to allow

For: ***(I) Non-factorizable models, (II) Models defined on the image lattice, (III) Starting point is unknown***.

- Our results are purely asymptotic

in N. They use Large Deviation Theory.

General Examples.

- Performance is determined by order parameters K as before:
- Example: K =1.00 and K =-0.43

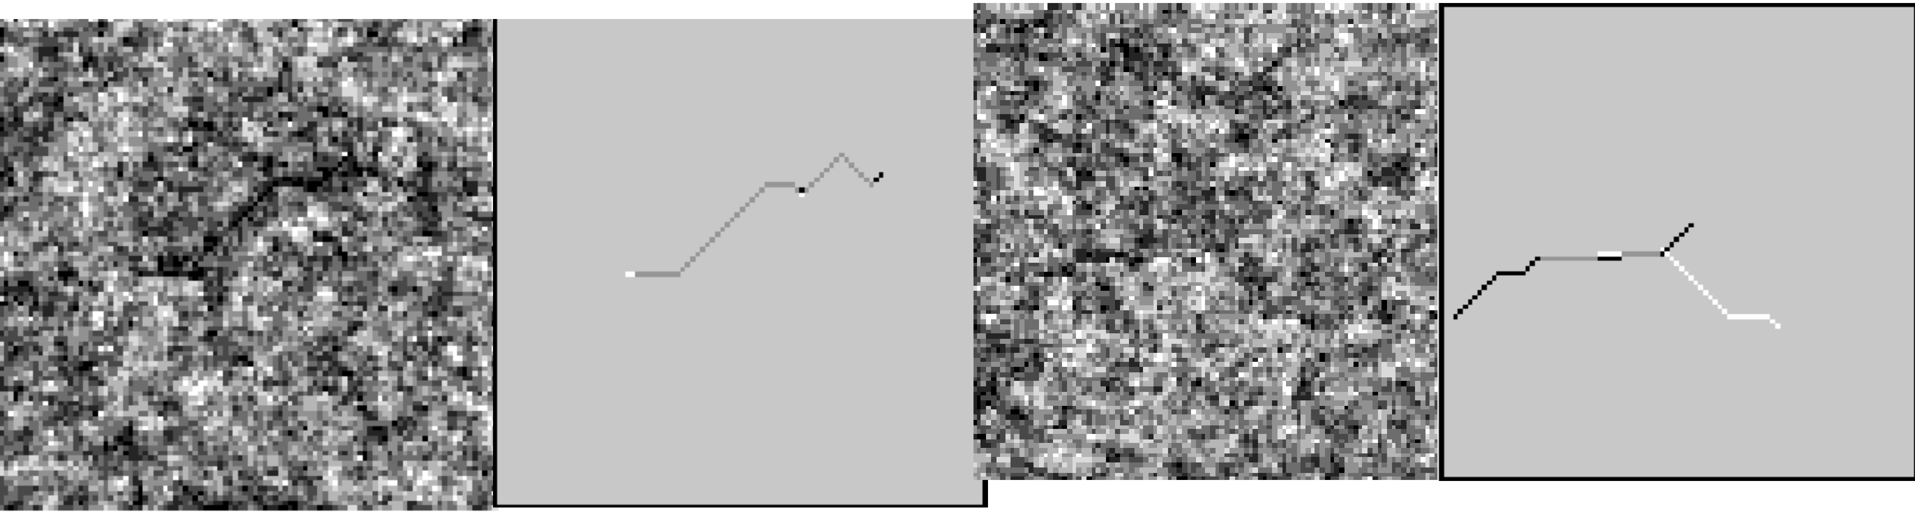

Approximate Models

- Similar Analysis for approximate models. (Sufficient statistics).
- Order parameters get reduced by terms such as $D(P_g || P_g^a)$.
- Requires the approximation P_g^a to use a subset of the statistics used by P_g See Minimax Entropy (Zhu, Wu,Mumford)

Summary

- Performance Bounds required for evaluating vision algorithms.
- Theoretical bounds can be determined in some cases. Order parameters K for characterizing task difficulty.
- Alternatively, frequentist measures on datasets with ground truth.